

The diagram illustrates a client-server system architecture. On the left, a **client PC** (102) is connected to an **Internet** (100). The Internet is represented by a series of horizontal lines showing the flow of data packets. The flow starts with a **SYN(connect)** packet from the client to the server, followed by a **SYN-ACK** packet from the server to the client, and then a **SYN(get)** packet from the client to the server. This is followed by a series of **file block** transfers from the server to the client, and then a **FIN (close)** packet from the client to the server, followed by an **ACK** packet from the server to the client, and finally a **FIN** packet from the client to the server, followed by an **ACK** packet from the server to the client.

On the right, the **server** (104) is shown. It consists of a **TCP listen queue** (106) and an **HTTP listen queue** (108). The TCP listen queue (106) receives connections from the Internet and passes them to the HTTP listen queue (108) via an **accept** operation. The HTTP listen queue (108) then passes requests to the **httpd daemon + server threads** (110). The httpd daemon (110) is represented by a vertical line with circles, indicating a queue of threads. The threads can **write** data to the **output buffers** (114) or **read** data from the **DISK** (112). The output buffers (114) are represented by a series of horizontal lines, and the disk (112) is represented by a cylinder. The **I/O controller** (116) is shown as a circle connected to the output buffers (114) and the disk (112).

FIG. 1

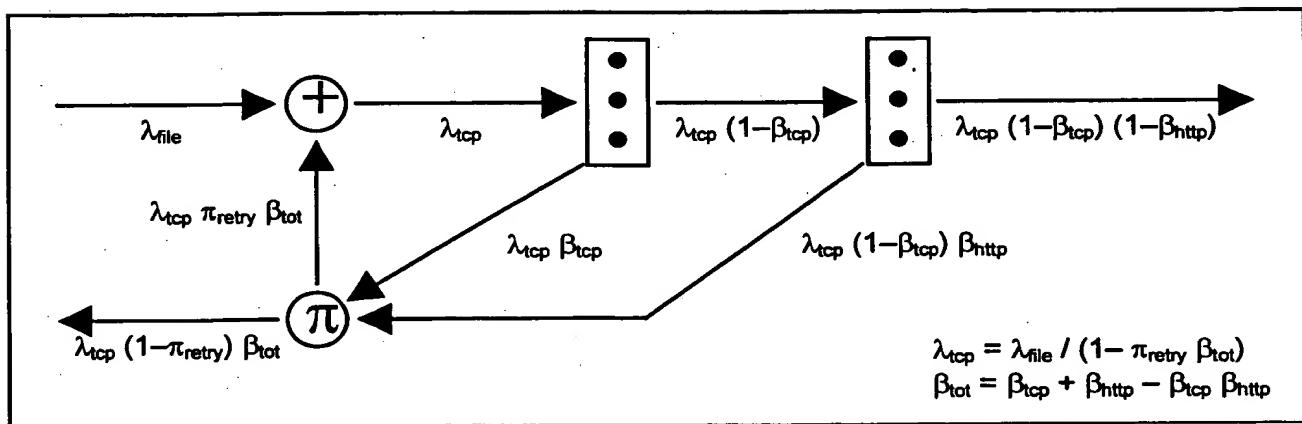


FIG. 2

Figure 3.a. Simulated and approximated server throughput versus transaction request rate

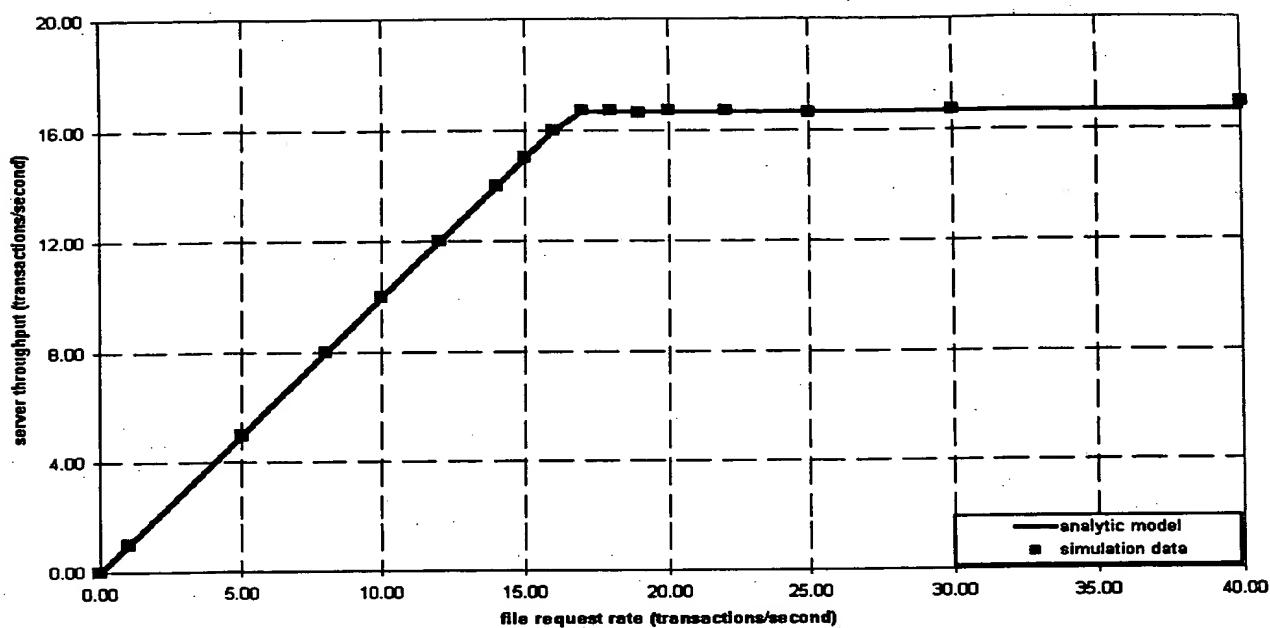


Figure 3.b. Simulated and approximated end-to-end service time vs. transaction request rate

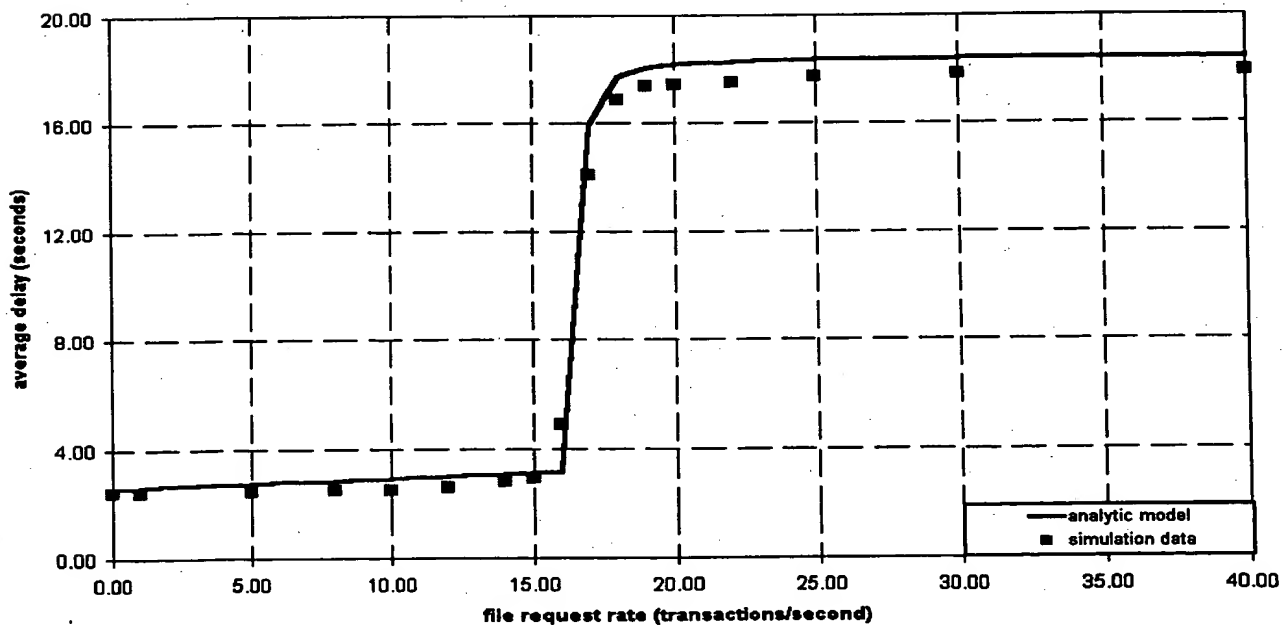


Figure 4.a. Simulated and approximated server throughput versus transaction request rate

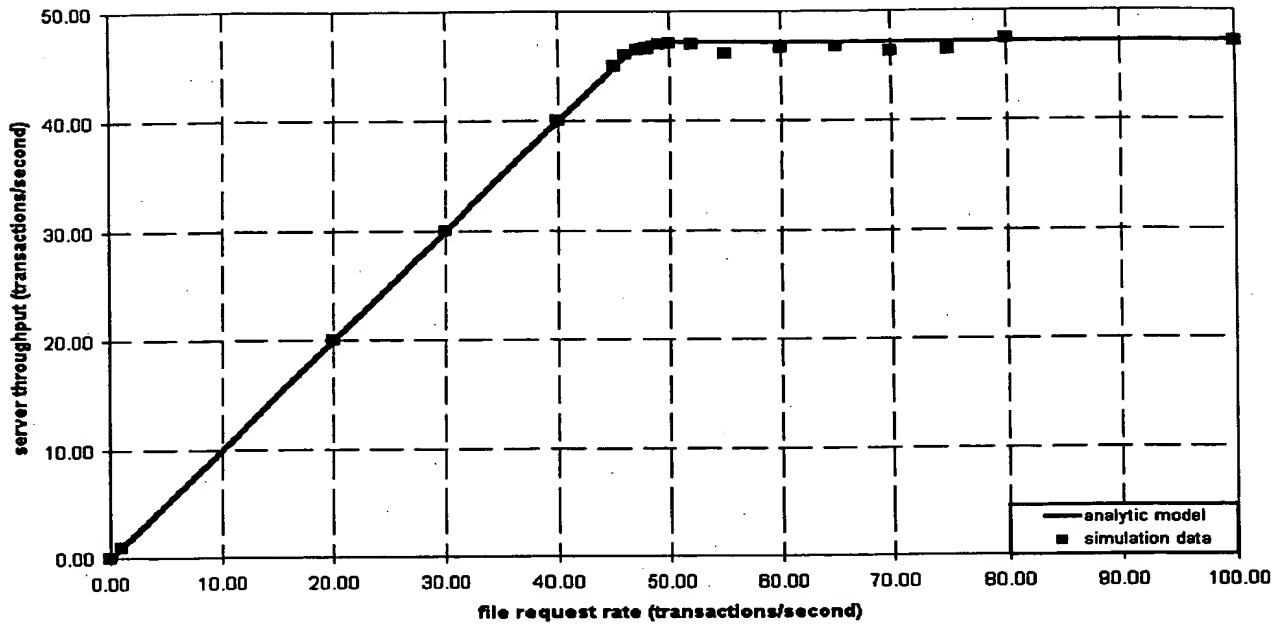


Figure 4.b. Simulated and approximated end-to-end service time vs. transaction request rate

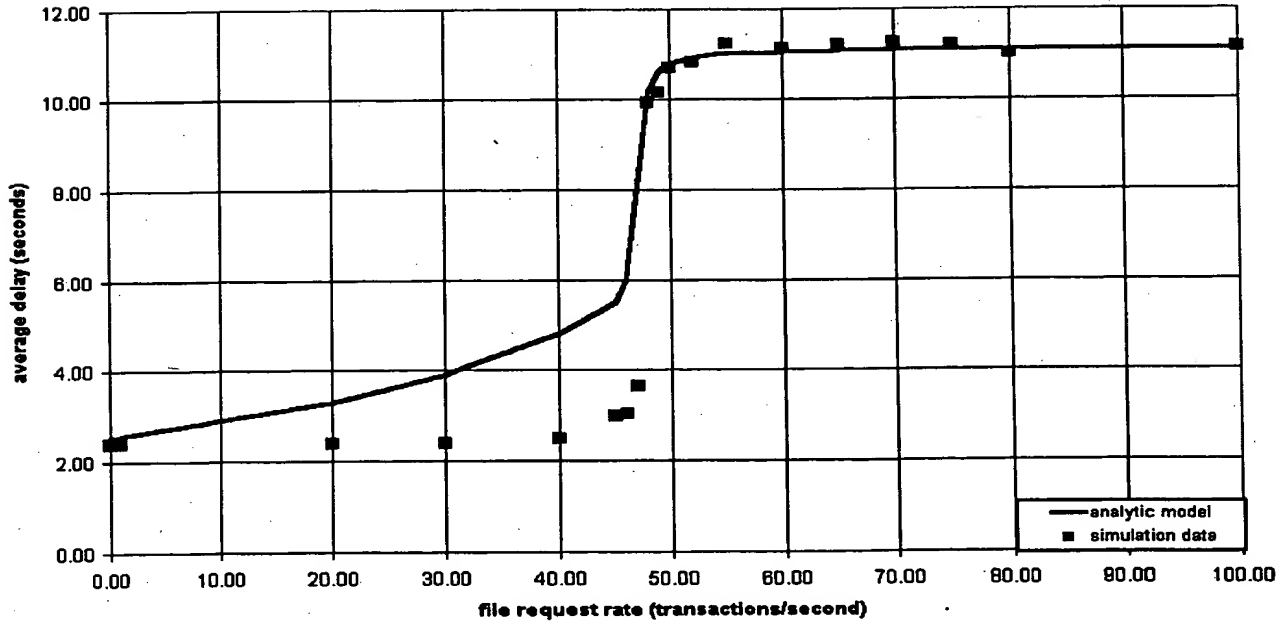


Figure 5.a. Simulated and approximated server throughput versus transaction request rate

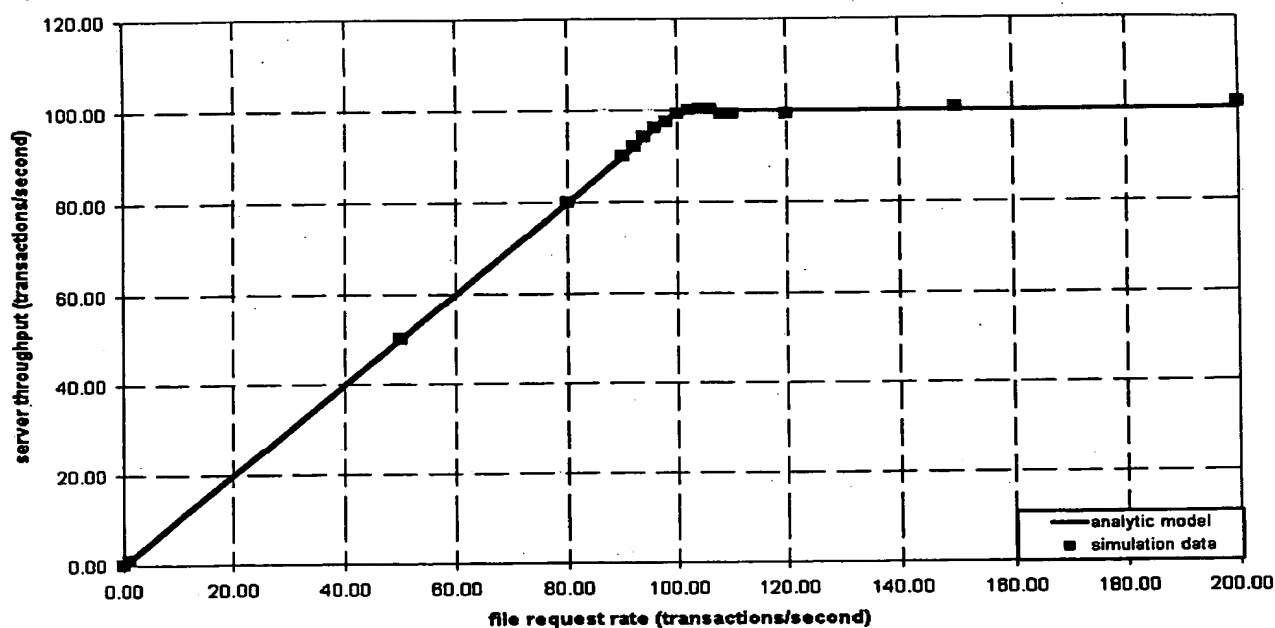
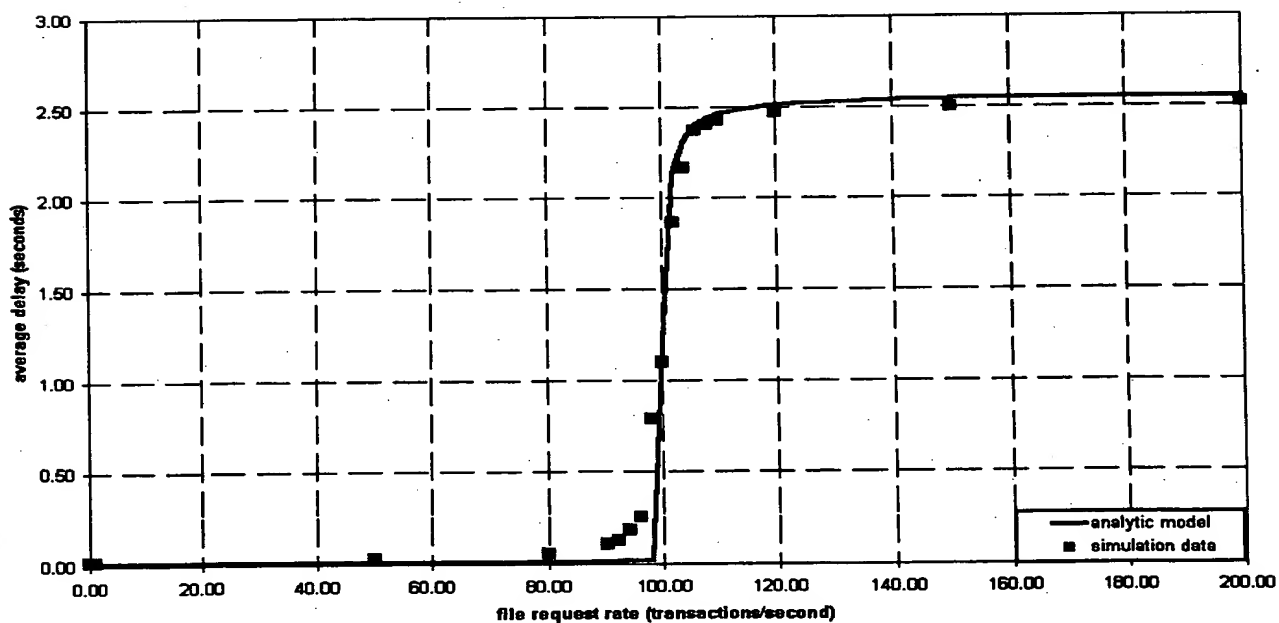


Figure 5.b. Simulated and approximated end-to-end service time vs. transaction request rate



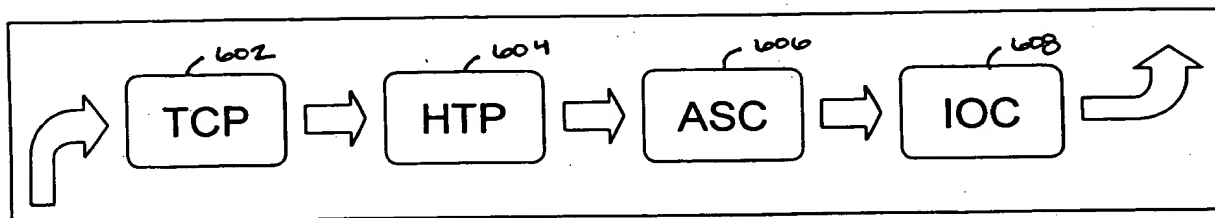


FIG. 6

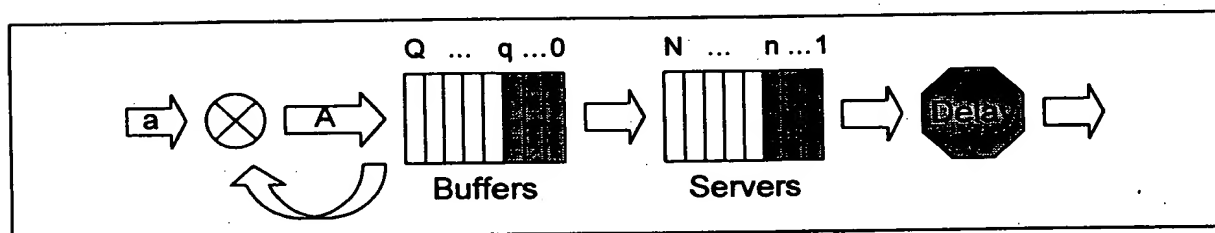


FIG. 7

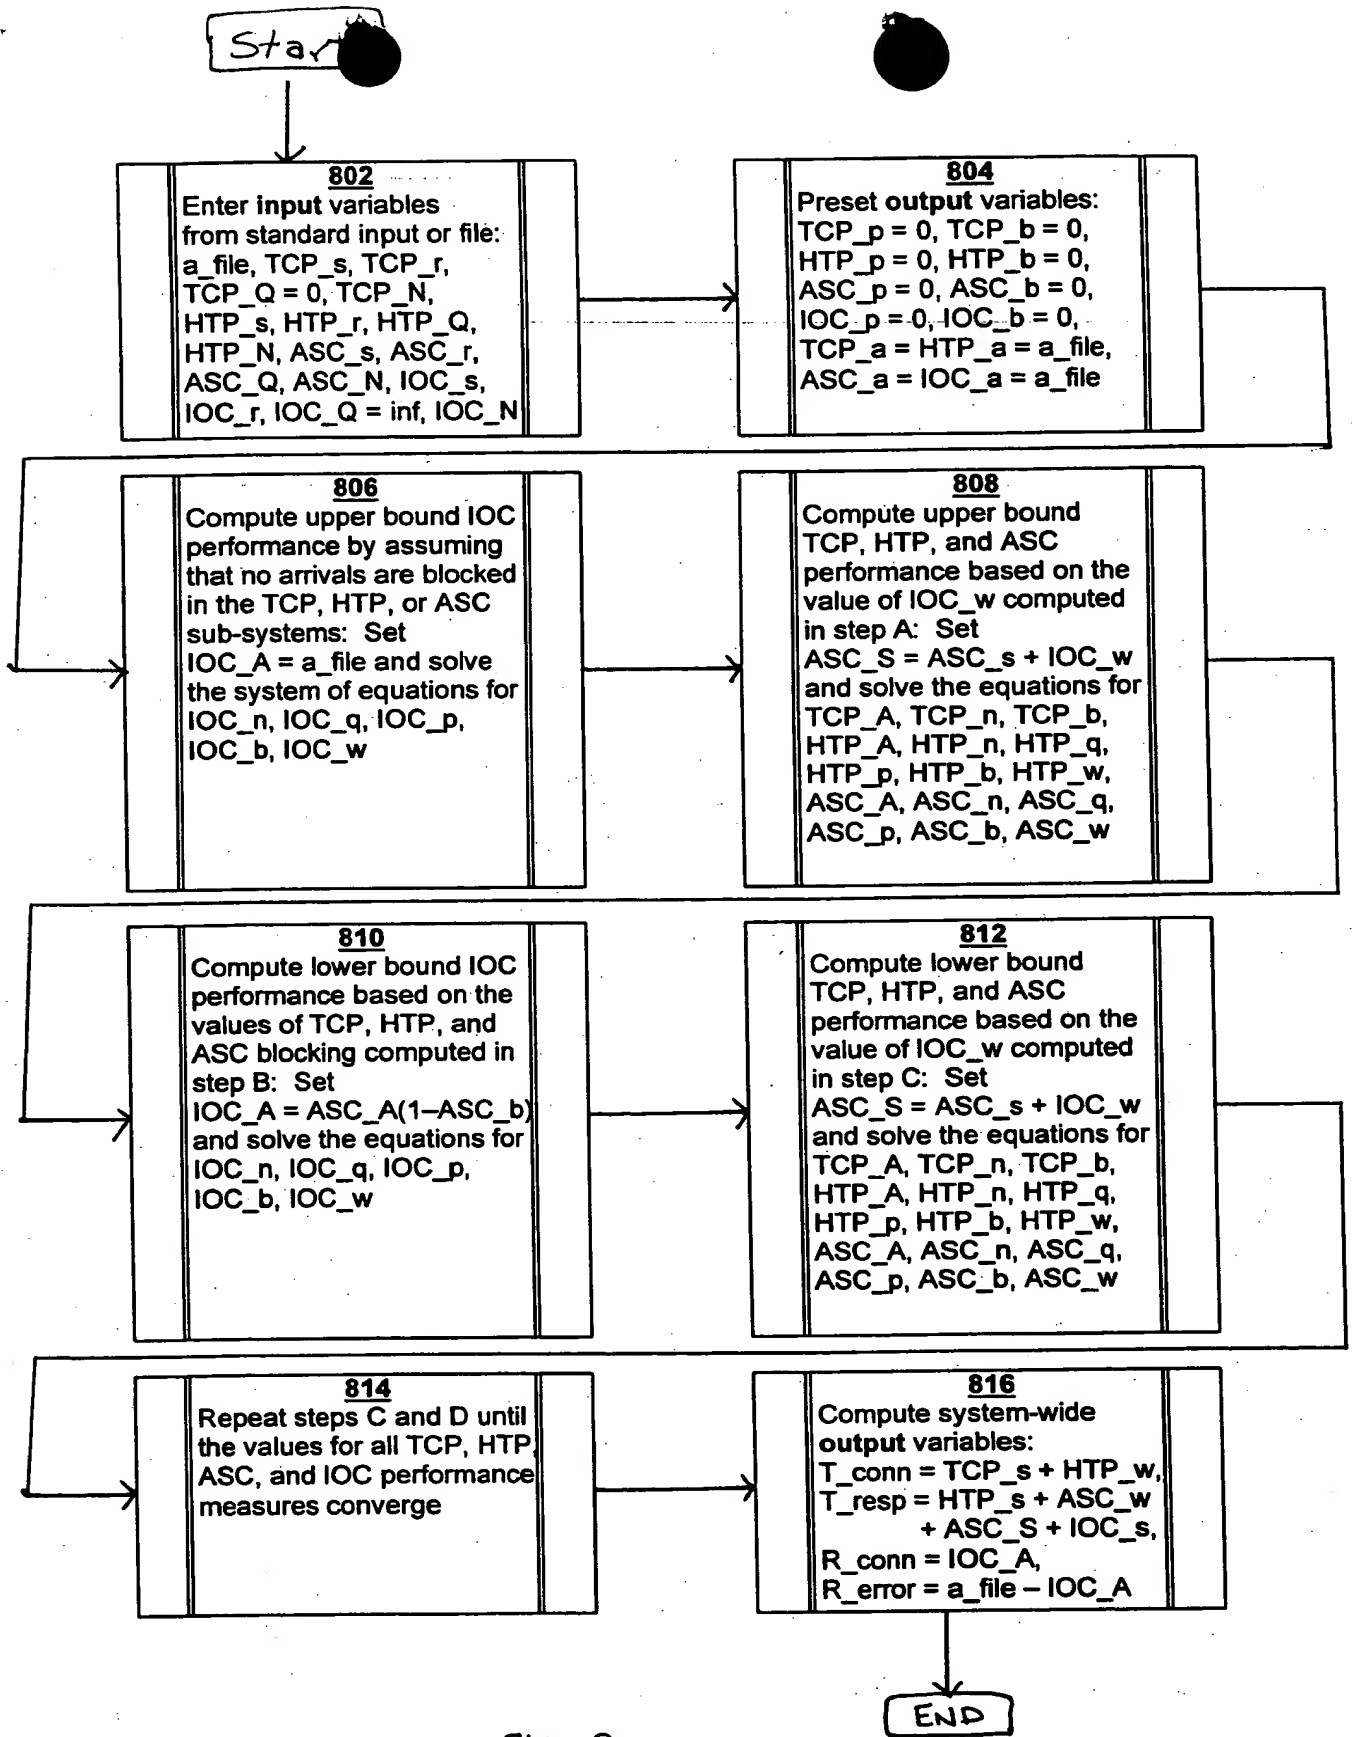


Fig. 8